



Need an Archive? But Want More Than Images?

**Evercore™ Clinical Information Manager
Solution Overview**

TeraMedica Product Development Services

*Timothy O'Connor
Vice President, Product Development*

*Damien Evans
Lead Architect*

TeraMedica Inc.
10400 Innovation Dr.
Suite 200
Milwaukee, WI 53226
866.290.8800
www.teramedica.com

Published February, 2007



CONTENTS

EVERCORE™ CLINICAL INFORMATION MANAGER	0
SOLUTION OVERVIEW	0
CONTENTS	2
INTRODUCTION	1
Problem Space.....	1
Solution Overview	1
ARCHITECTURE	2
Overview	2
Enterprise Java	3
CLINICAL CONTENT	5
DICOM Content.....	5
Non-DICOM Content	7
HL7 Content	10
Multiple Patient Medical Record Number Support.....	11
SMARTSTORE	12
Overview	12
Organizational Hierarchy	12
Storage Policies	14
Lifecycle Management Policies	17
Pediatric Policies	17
UNIVISION - EMR/EHR INTEGRATION.....	18
Overview	18
Evercore UniVision	18
SECURITY	20
Overview	20
HIPAA.....	20
SmartStore Security	20
Audit Trails.....	20
Enterprise Security Integration	21
DEPLOYMENT	22
Overview	22
Scalability	22
Deployment Options	22
FastGrid: Evercore Distributed Deployment	23
GLOSSARY	26

The information contained in this document represents the current view of TeraMedica Inc. on the issues discussed as of the date of publication. Because TeraMedica must respond to changing market conditions, this document should not be interpreted to include commitments on the part of TeraMedica Inc., and TeraMedica cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. TeraMedica, Inc. MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION CONTAINED IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of TeraMedica Inc.

TeraMedica may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from TeraMedica Inc. The furnishing of this document does not provide for any license to these patents, trademarks, copyrights, or other intellectual property.

© 2001 to 2007 TeraMedica Inc. All rights reserved.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

INTRODUCTION

Valuable pieces of the longitudinal patient record are being locked away in proprietary archives, stored unsecured in log books or folders, and generally out of reach to those responsible for patient care.

Problem Space

From film libraries to PACS, clinical content in healthcare has traditionally been managed by the department that created it. Clinical areas are often charged with providing solutions to maintain their data for near-term treatment and diagnosis as well as long-term storage and research. Valuable pieces of the longitudinal patient record are being locked away in proprietary archives, stored unsecured in log books or folders, and generally out of reach to those responsible for patient care.

Chief executives of healthcare providers need a solution to unlock this clinical information and provide access to the data. The digital revolution in clinical information, along with regulatory and cost concerns is forcing a paradigm shift in the way providers both value and maintain their clinical data assets.

Solution Overview

The TeraMedica Evercore system creates a means to bring together pieces of clinical content under one institutional infrastructure. It accomplishes this while providing unique qualities of service to the originating clinical areas and relieving them of the burden that accompanies long-term management of critical patient information.

Physicians can procure and utilize the most care-effective departmental workflow and visualization products while knowing that their clinical data is managed according to their unique policies and rules in a secure, robust, IT managed solution.

Chief Information Officers can provide unique services while maintaining the disparate clinical assets under a common patient record to facilitate sharing the data (to an Electronic Medical Record or Electronic Health Record, for example). This is all accomplished in a robust enterprise infrastructure that conforms to corporate IT standards and practices for data security, availability, backup, etc.

This document provides insight into the unique features, functionality, and benefits of the TeraMedica Evercore solution.

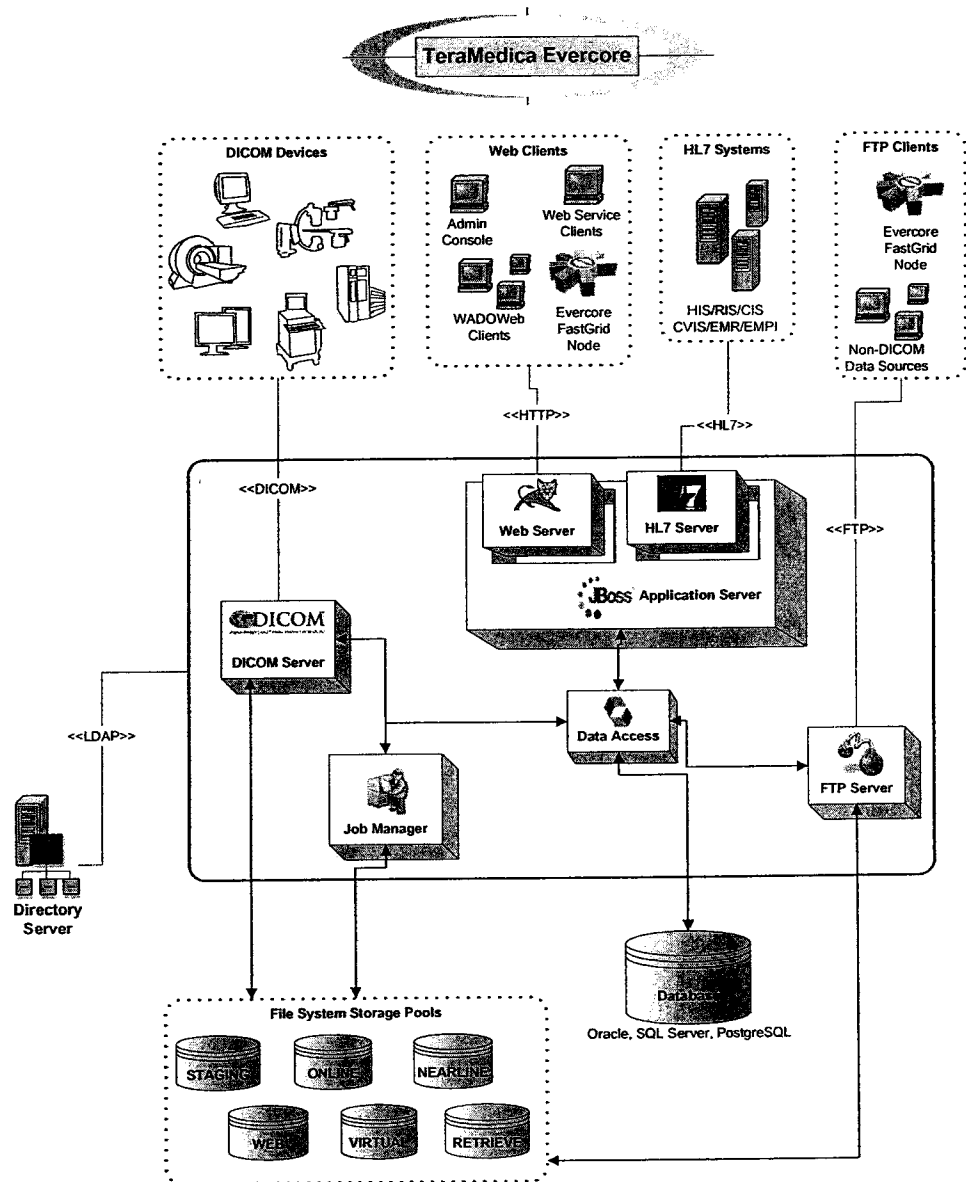
Evercore Clinical Information Manager brings together all clinical content under one common infrastructure.

ARCHITECTURE

With no restrictions on technologies or methodologies... a unique blending of an enterprise-scalable architecture with healthcare-domain capabilities.


Overview

TeraMedica was fortunate to be able to start development of the solution with no restrictions on technologies or methodologies. This freedom allowed for a unique blending of an enterprise-scalable architecture with healthcare-domain capabilities. The architecture enables Evercore to fit into an IT infrastructure at either the departmental or enterprise level. The system features a service-oriented set of medical domain (DICOM, HL7) and IT (HTTP, XML, RMI) interfaces, backed by components which perform real time transaction and offline job processing.



Enterprise Java

The Evercore solution has been primarily developed using the Java programming language. Java was selected as the development platform due to its suitability for server-side applications, portability, and performance. In addition to the standard Java development toolkit, Evercore utilizes several portions of Java Enterprise Edition (JEE). JEE standards such as guaranteed messaging, remote service access, and resource pooling are essential for a critical enterprise information system. The C programming language is used within the image compression logic in order to gain the most efficiency when performing resource intensive operations.



Each of these modules can be individually tuned, replicated, and configured as needed in order to satisfy installation requirements.

Modularity

Evercore separates transactional areas of functionality into modules, some of these modules having their own software processes. These areas include DICOM, HL7, Web, and Jobs (asynchronous, offline, and/or batch operations). Transactions and resource pooling occur within. Evercore itself (as an entire application) can be replicated on a multi-homed host machine in order to take advantage of available resources. In cases where multiple servers are running Evercore, a load balancing switch can be utilized to distribute requests across the servers.


The modularity of the software and the wide flexibility of configurations allows for Evercore deployments that scale to meet the needs of the largest medical imaging environments, while still being able to service only a single department if necessary. This modularity lends itself well to the ability to fail over if components do not function properly. Replication of components and replication of Evercore instances (on a single server or multiple servers) provides a high degree of assurance that business continuance can be achieved.

Job Management


The job manager processes are responsible for coordinating and executing all offline and batch operations. This is due to the fact that some operations are well suited to an offline scheduled procedure while other operations can benefit from being run asynchronously without forcing a user to wait for the end result. The job management infrastructure is a robust implementation which allows for scheduling, retries, pre and post condition workflow processing, thread pooling, and job execution policies (defining the runtime characteristics of each job type).

Job managers maintain several thread pools that allow for different job execution policies. These policies define the maximum number of concurrently executing jobs, what to do if the pool is full, number of retries, etc. Jobs are guaranteed to be updated at the end of processing, regardless of the outcome, and have the ability to notify an administrator via email if there is a failure. A job manager is capable of checkpoint restarting in cases of premature termination and also allows for advanced workflow through the use of pre and post conditions.

A common deployment scenario involving job management entails multiple job managers per server handling different job types throughout the day, but all working together at night to perform archiving to long term storage.



Some operations are well suited to an offline scheduled procedure while others can benefit from being run asynchronously without forcing a user to wait for the end result.



A specialized TeraMedica-developed LDAP schema stores variables and policies in an organizational hierarchy.

LDAP Rules and Configuration

LDAP (Lightweight Directory Access Protocol) is utilized to store the bulk of the system configuration. A specialized TeraMedica-developed LDAP schema stores variables and policies in an organizational hierarchy. This enables the configurations and policies to be applied to specific business units that are mapped (configured) at specific points in the organizational tree.

A simple example would be a DICOM device being added under the Radiology organization node. The DICOM device would have the policies and configuration relevant to that Radiology department applied to all of the operations stemming from that device's interactions with Evercore. Evercore processes register with LDAP to receive update notifications when the configuration or policies change. This allows the process to dynamically adapt to the changes without restart.

Utilizing LDAP for this purpose provides for a fast, dynamic means of synchronizing configuration across an Evercore cluster. It also has the added benefit of being able to leverage an existing Directory Server to obtain user authentication and authorization information.

CLINICAL CONTENT

DICOM Content

Overview

Evercore adheres to the Digital Imaging and Communications in Medicine (DICOM) industry standard for file format and communication to imaging modalities, PACS, and other compliant systems or workstations. The Evercore system acts as both a DICOM Service Class User (SCU) and a DICOM Service Class Provider (SCP) in certain instances as outlined in the *Evercore DICOM Conformance Statement*. Please consult the *Evercore DICOM Conformance Statement*, available at www.teramedica.com for more detailed information regarding how DICOM dataset and transport protocols are used within the Evercore solution.

DICOM Toolkit

TeraMedica maintains our own DICOM development toolkit. This provides great flexibility and timeliness in supporting new DICOM features and functionality without relying on a third party. The toolkit is written entirely in Java and is proven and stable in high transaction environments.

DICOM Basics

Object Storage

In order to accept incoming DICOM data objects for storage and management, Evercore acts as a Service Class Provider (SCP) for the DICOM C-STORE command. Configuration of DICOM study close policies can be made at the device/connection level and can be association based, timer based or storage commitment based.


At the time of storage, certain DICOM header elements are read from the DICOM object and stored in the Evercore database. For all incoming DICOM objects, configurable business policies dictate how the studies are stored. These business policies indicate how many copies of the study to store, what data compression types and ratios to apply, the data retention period, ownership of the data, etc. (see SmartStore). The incoming objects can optionally be validated against the DICOM standard to insure appropriate compliance.

Query/Retrieve


To provide for the query and retrieve of stored studies, series or images, Evercore acts as a Service Class Provider (SCP) for both the DICOM C-FIND command and the DICOM C-MOVE command. A variety of DICOM data elements can be used for query purposes and returned studies, series, or images can (optionally) have the DICOM header data coerced with up-to-date database fields. As part of the DICOM C-MOVE command, Evercore also acts as a DICOM Storage Class User (SCU) for the C-STORE command to send the requested DICOM objects to the designated DICOM target(s). The format for the returned DICOM objects is based on configurable business rules.

Query Spanning

In addition to providing for standard DICOM query requests for data stored in its own system, Evercore supports the spanning of those queries to other DICOM systems. A single DICOM query to an Evercore system can



For all incoming DICOM objects, configurable business policies dictate how the studies are stored.



A variety of DICOM data elements can be used for query purposes. Returned studies, series, or images can have the DICOM header data coerced with up-to-date database fields.

automatically generate additional DICOM queries to one or more additional DICOM C-FIND service class providers. The results of these queries are then coalesced into a single response that is passed back to the query client. The Client can then retrieve the data it wishes directly from the system(s) where it is stored.

Storage Commitment

Evercore supports the push model of DICOM Storage Commitment as a SCP and responds to the DICOM N-ACTION command. The Evercore system can be configured to respond positively to a DICOM storage commit query only when certain criteria are met related to the associated objects study status (i.e. commit with no Quality Assurance (QA) issues, commit when stored on near-line storage, etc.).

Export

The Evercore web browser interface provides the ability for a user to send (as a C-STORE SCU) DICOM studies and/or series from Evercore to one or more DICOM targets. In addition, Evercore provides the ability for users to view real-time results of their DICOM transfers.

Verification

As required by the DICOM standard, Evercore supports DICOM verification C-ECHO command as both an SCU and SCP.

Object Validation

To insure that only valid DICOM data is accepted, the system can be configured to validate the DICOM conformance of incoming DICOM objects. This feature can be enabled as a storage policy to act only on a subset of the data stored into Evercore and is useful for verifying other vendors DICOM compliance.

Tag Mapping


The Evercore system provides a DICOM field mapping capability which allows many poorly behaved DICOM devices to store data into Evercore. For these devices, DICOM fields within the incoming data can be mapped to Evercore database fields. Each device configured within Evercore can have a unique field mapping template.

Transformations

In addition to Tag Mapping, Evercore provides the ability to perform transformations on the incoming DICOM data prior to population of the database. A default set of transformations are provided with the system and additional transformations can be added by including the appropriate Java classes. This powerful functionality can be used for a variety of purposes including complex manipulation of the data, providing for default values, validation against other files or tables, etc. These transformations can be configured to be applied only to specific DICOM devices if so desired.

Data Coercion

The system can be configured to optionally coerce all changes to DICOM related fields into the DICOM header upon export from Evercore. This means that when changes are made to patient and/or study metadata in the Evercore database (either manually or via an HL7 interface), those changes can be reflected in the DICOM header when content is sent out of Evercore. The fields



to be coerced can be configured and can vary by organization. If so configured, the original DICOM dataset is always stored unmodified on the file system and available for medical-legal purposes.

Auto-forwarding

Evercore's auto-forward functionality provides the capability that as new studies are accepted into the system, a copy can also be sent via DICOM to one or more DICOM targets (specialty workstation, PACS, etc.). This feature can be enabled as a storage policy to act only on a subset of the data stored into Evercore.

Policies related to the application of compression and the number of copies to maintain are configurable at the organization level providing business-unit unique qualities of service.

Image Compression

Evercore supports a variety of medical compression formats for storing incoming data. In addition to storing the data as received, Evercore can be configured to apply modality specific compression to the data. Multiple copies of the studies can be maintained in the Evercore system in different compression formats if so desired. Policies related to the application of compression and the number of copies to maintain is configurable at the organization level providing business-unit unique qualities of service. Among the compression formats supported are:

- RAW
- DICOM JPEG (lossless and lossy)
- DICOM JPEG2000 (lossless and lossy)
- JPEG (lossy – for lightweight web distribution)
- JPEG2000 (lossy – for lightweight web distribution)

Performance Monitoring

As part of its system monitoring capabilities, Evercore includes the ability to view DICOM association performance for devices with which it interacts. The association results can be filtered by a variety of filters including; called AE Title, calling AE Title; destination AE Title; host; DICOM command; success status; and date range. The individual DICOM commands within each association can also be viewed with query/modality code; study series and image counts; best and worst performance; bytes transferred, destination AE Title, and time to complete.

Non-DICOM Content

Overview

There exists a wide variety of clinical data which does not fit well into the DICOM standard as well as many departmental information systems that have not embraced the standard. In clinical areas that have been traditionally underserved by PACS vendors, patient digital data is often poorly maintained and inaccessible to the wider enterprise. The Evercore solution provides a foundation for the management of non-DICOM clinical assets using institutional Information Technology standards and practices. In addition to ensuring that valuable patient data is maintained in a managed environment, Evercore can make this data more readily available throughout the broader enterprise.

The Evercore solution provides a foundation for the management of non-DICOM clinical content using institutional Information Technology standards and practices.

The non-DICOM content is retained within Evercore as it was received and when requested, is provided in the format in which it was stored.

The Basics

Virtually any type of digital clinical content that is related to a patient or a patient's study can be stored and maintained within the Evercore system. This content is organized using a convenient folder-based structure which provides for multiple folders per patient and/or a patient's study. Folder descriptions allow for easy identification of contained content. In addition, the system can be configured to accept only certain types of content, thus preventing users from storing unauthorized data.

After the non-DICOM content is stored within Evercore, policies much like those available for DICOM content can be applied for the storage and management of the data. The non-DICOM content is retained within Evercore in its native format (as it was received) and when requested, is provided in the format in which it was stored.

Labels (keywords)

In order to provide for the classification of content within Evercore, one or more labels can be associated with the content. These labels can also be used as search criteria when looking for specific types of stored content, and new labels can be added into the system or to content as necessary. The labels available for users to attach to content can be restricted based on the user's organization and authority.

Import Connectors

To facilitate the storage of non-DICOM content, Evercore uses a variety of import "connectors" which allow users or external systems to interact and store content in the system. The modular design utilized for import allows TeraMedica to add additional connectors in the future. The current set of non-DICOM import connectors available include:

Only users authorized to access non-DICOM content, see and interact with the content using the Evercore user interface.

- **User Interface**

Through the user interface, users can easily upload and attach content either directly to a patient or to particular studies for a patient. Multiple pieces of content can be uploaded at one time and labels can be associated at upload time. Content is available for users to download immediately upon successful upload.

Access to non-DICOM content is a user-based authorization. That is, only users authorized to access non-DICOM content see, and are able to interact with the content using the Evercore user interface.

- **Shared Directories/FTP**

Evercore can be configured to actively monitor one of more shared directories for incoming clinical content. Data can be placed directly into these directories or received via an FTP interface. A descriptive file must be provided with the clinical content which provides details on the patient or patient/study context for the content to be associated with, along with identifying descriptive metadata.

- **Web-Services**

A set of documented web-based services interfaces are available which allow external systems to integrate with Evercore for the exchange of clinical content. These interfaces provide for upload, query, and retrieval of non-DICOM content.

Export Connectors

In addition to the import connectors, Evercore provides export connectors for retrieving or exporting non-DICOM content out of the Evercore system. In all forms of export, the non-DICOM content is returned to the user in the exact same format as it was originally stored into the Evercore system. The current set of non-DICOM export connectors available include:

- ***Direct Download via User Interface***

Users can download non-DICOM content directly via the Evercore web interface. The content can be saved to a file or opened automatically using the application configured by the web browser (i.e. Adobe Acrobat Reader for .PDF files).

- ***XDS Document Source***

The Evercore system can act as an XDS (Cross-Enterprise Document Sharing) Document Source for exporting non-DICOM content to an XDS Registry/Repository. The user can select an XDS target from the user interface to export the data to.

- ***Shared Directory/FTP***

From the Evercore web interface, users can export non-DICOM content to a shared directory or to an FTP server. An FTP user ID and password can be specified by the user for FTP target authentication. The status of the exports can be monitored by the user via the web interface.

- ***Web-Services***

A set of documented web-based services interfaces are available which allow external systems to integrate with Evercore for the exchange of clinical content. These interfaces provide for upload, query, and retrieval of non-DICOM content.

The unique organizational structure of Evercore provides support for multiple departmental or enterprise interfaces, with each applicable to a logical segment of the entire data stored within Evercore.

HL7 Content

Overview

Evercore provides unique mechanisms to facilitate the interfaces required by a variety of information systems found throughout healthcare provider networks. Evercore uses the Health Level 7 (HL7) standard to interface with various clinical, departmental or enterprise information systems including (but not limited to) Radiology Information Systems (RIS), Cardiovascular Information Systems (CVIS), Hospital Information Systems (HIS), and Master Patient Index (MPI) systems.

In addition, the unique organizational structure of Evercore provides support for multiple departmental or enterprise interfaces, with each applicable to a logical segment of the entire data stored within Evercore.

The HL7 support allows Evercore to act as a "source" system for image study requests by ensuring that patient and study metadata is always kept up-to-date. This also allows Evercore to be the source system for serving all image requests from an enterprise EMR/EHR system.

HL7 Toolkit

TeraMedica maintains our own HL7 development toolkit. This provides great flexibility and timeliness in supporting new HL7 features and functionality without relying on a third party. The toolkit is written entirely in Java and is proven and stable in high transaction environments.

HL7 XSL Transforms

Through the use of XSL (Extensible Stylesheet Language) transforms, new HL7 message types and formats can easily be added to the system. The transforms can execute a series of defined actions within the Evercore system for each message received.

HL7 Messages

Study/Series Notification (outbound)

When datasets are stored into Evercore, information related to the studies and series can be sent to one or more HL7 systems (as a custom HL7 message format) as defined within an organizational structure. These interfaces are configurable at the Storage Policy level.

Study Validation (outbound)

Real-time HL7 (custom HL7 message format) validation of patient and/or study information can be verified at the time a study is stored. Incorrect information can be updated automatically or quality assurance (QA) issues can be generated for technologist review. These interfaces are configurable at the Storage Policy level.

HIS/CIS/MPI Patient Demographic Update (inbound)

Evercore can be configured with an HL7 interface (ADT HL7 message formats) to a Hospital Information System (HIS), Clinical Information System (CIS), Master Patient Index (MPI) system for patient demographic information updates (update, merge, etc.). This is configurable at the patient management organization level.


The HL7 support allows Evercore to act as a "source" system for image study requests by ensuring that patient and study metadata is always kept up-to-date.

Order Feed (inbound)

The system can be configured with an HL7 interface (IHE HL7 message format) for receipt of departmental imaging orders from a Radiology Information System (RIS), Cardiovascular Information Systems (CVIS), or other information system. These interfaces are configurable at the Storage Policy level.

Delete (inbound)

The system supports an HL7 inbound interface (custom HL7 message format) for requests to delete existing studies and series from Evercore. These requests typically are sent from departmental imaging or information systems. These interfaces are configurable at the Storage Policy level.



Multiple Patient Medical Record Number Support

It is possible (and not uncommon) for a patient to be assigned multiple Medical Record Numbers (MRN) in a healthcare provider network. Evercore is capable of storing multiple local MRNs assigned to a patient, and associating those with a single patient record in the system's database.

The system relies on the MRN type being configured at the Study Management Organization (STMO) level. Studies stored for a patient from within that STMO are associated with that MRN ID type. Evercore can utilize HL7 to communicate with a Master Patient Index (MPI) to resolve patients and maintain an external MPI identifier.

SMARTSTORE


Overview

Evercore has been designed and architected with a unique and highly adaptable method of modeling healthcare provider business units for use in the application of content management policies and the logical/physical segmentation of the clinical data.

This organizational context permeates much of the functionality that defines the Evercore solution and can provide distinct qualities of service to clinical areas that utilize the system.

Organizational Hierarchy

The organizational hierarchy feature of Evercore can provide healthcare systems with a shared enterprise infrastructure for managing clinical content, while maintaining flexible business policies implemented to meet the needs of individual hospitals and/or departments.

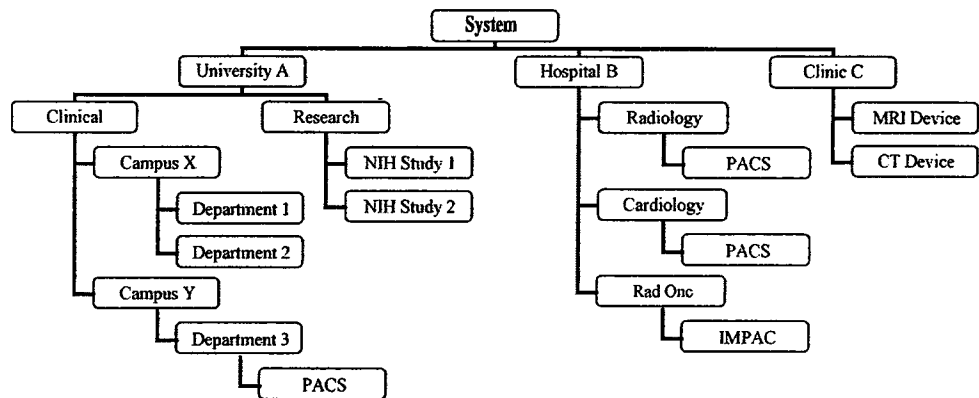


A hierarchical representation of nodes, with parent-child relationships, represents each business unit as defined by the deployment requirements. This organizational mapping allows representation of organization structure from the very simple (flat) to the complex (nested trees). Evercore can be configured to support multiple organization levels. The hierarchical tree can be defined to map to the layout of the deployed provider and change as the organization changes through merger, acquisition, or consolidation.

The system can be configured to manage the data and transactions for multiple sites, facilities, departments, and business units within the same shared infrastructure. These data and transactions are managed uniquely, each with their own quality of service, for that organization. Organization nodes are specified for patient, study and device ownership and provide a high degree of configurability for:

- Defined Patient Management Organization (PMO) levels for patient ownership and maintenance
- Defined Study Management Organization (STMO) levels for study ownership and maintenance
- Application of storage policies by organization node
- Integration to external systems via HL7 by organization node
- Enhanced granularity and application of data security (need-to-know access)

Organizational mapping is a way of grouping and identifying subsets of patient and study content for the purpose of applying data management clinical policies. This provides the ability for a healthcare provider to have one common data storage and management infrastructure, but have data from their clinical business units managed uniquely and independently of one another.



Organization Hierarchy (example)

In most cases an organization node takes on the configuration and/or policies of its nearest parent. When a device stores data it stores to the storage policies defined in the first parent organization with storage policies defined. The first parent organization is found by traversing up the tree, beginning at the client's organization node.

There are two special types of organization nodes; Patient Management Organizations (PMO) and Study Management Organizations (STMO).

- **PMO:** A PMO node denotes an organization level that is responsible for ownership of patient records (as known by the Evercore system). A patient MRN must be unique within a given PMO.
- **STMO:** This node denotes an organization level that is responsible for the ownership and management of clinical study records.

Content Mapping enables Evercore to operate with available storage hardware and/or software and allows for the dynamic addition, removal, and migration of physical storage within the system.


Storage Policies

Content Mapping

Content Mapping is a term which is used to encompass the proprietary technology Evercore uses to abstract, or “virtualize” physical storage. Content Mapping enables Evercore to operate with almost any available storage hardware and/or software and allows for the dynamic addition, removal, and migration of physical storage within the system. This technology is comprised internally of Storage Policies and Storage Pools.

Storage Policies and Storage Pools

A Storage Pool represents a logical view of the physical storage resources and provides the means for storing to those resources through organization specific Storage Policies. Storage policies are defined within the context of a Storage Pool, and provide properties and rules for the content according to that policy. When data is stored into Evercore, it is assigned to existing storage policies, and stored to the Storage Pools to which those policies belong.



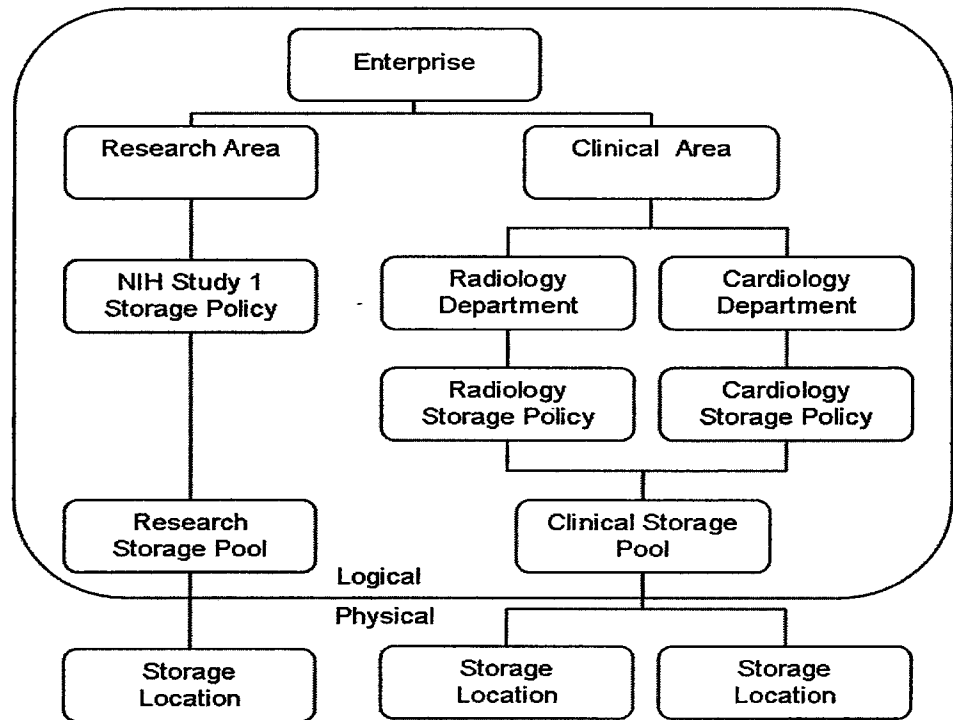
Storage Policy Flow (simple version)

The Storage Policies determine what content are stored in its Storage Pool and how that content is handled (in relation to compression, retention periods, validation, notifications, auto-forwarding, etc). As part of SmartStore, the Storage Policies define the storage behavior for a given organization and a given set of content.

The Storage Pool is responsible for coordinating the activities between the Storage Policies and the physical media. The Storage Pool is a logical container for one or more physical locations on disk/tape/etc, providing access to those locations in the manner defined by the policies in the pool. A Storage Pool can be thought of in a similar manner as a database connection pool; instead of managing database connection resources; the Storage Pool manages storage resources.

The logical Storage Pool is responsible for coordinating the activities between the Storage Policies and the physical media.

Storage locations are a representation of a physical storage entity (directory, drive, mount point, file system, etc.) which are stored to only by the Storage Pool to which it belongs. Storage locations can reside on direct attached storage (DAS), storage area networks (SAN), network attached storage (NAS), and grid storage, etc.



SmartStore - Storage Relationships to Organizations (example)


Storage Policies define the storage behavior for a given organization and a given set of content.

Types of Storage Policies

There are several types of storage policies and although this section only refers to storage policies, their parent storage pools also fall into the same classifications:

- **Staging Policy:** The staging storage policy is a system reserved area where raw objects are initially placed by the system. Storage dispatcher and archive processes work with data in this staging area.
- **Online Policy:** An online storage policy is for content typically stored to spinning disk media (i.e. RAID). DICOM content stored using online storage policies are maintained in DICOM Part 10 format according to the modality-specific compression specified.
- **Nearline Policy:** A nearline storage policy is typically used for archival to long-term media (i.e. tape DVD, etc.). DICOM content stored to a nearline policy is placed into a single file, one file per study. Data retrieved from a nearline policy gets extracted to an associated retrieve policy.
- **Retrieve Policy:** A retrieve storage policy is a type of online policy that is only utilized when data is moved back from a nearline policy. By not using a standard online policy to move data to, the online policy is not made "dirty" with older de-archived data. The nearline single data file is extracted to the retrieve policy and the retrieve policy rules are applied to the data. Typically a retrieve policy would have a shorter retention period than a standard online policy.
- **Web Policy:** A web storage policy is a type of online policy which maintains highly compressed JPEG images for use by a web-based distribution interface (such as Evercore UniVision).

- **Virtual Policy:** Virtual storage policies are short term online policies where content resides while being transferred to a centralized system using the Evercore FastGrid module at remote distributed nodes.



Storage Policy and Storage Pool Flushing

There are two types of storage flushing available in Evercore SmartStore. Retention-based flushing occurs at the Storage Policy level and is based on the retention rules set for that policy. Critical flushing happens at the Storage Pool level and also utilizes rules when it needs to flush data. The major difference between these two types of flushing is that retention flushing happens proactively for an organization's storage due of their data management policies, while critical flushing happens reactively in the system because of poorly allocated storage for a given Storage Pool.

Storage Policies trigger their own retention-based flushing based on defined business rules. This retention-based flushing is accomplished by a scheduled job within the system.


The following flushing rules are available:

- **Least Recently Accessed:** Content is flushed only if they have been accessed (retrieved) prior to a configurable time period. This is appropriate for a retrieve Storage Policy.
- **Least Recently Created:** Content is flushed only if it was created in the Evercore system prior to a configurable time period.
- **Least Recently Performed:** Content is flushed only if it has a performed date prior to a configurable time period.

Critical flushing of Storage Pools occurs if storage space is exhausted. This is a situation that *should* never occur if the system has been configured properly. When a Storage Pool enters critical flush mode, it forces its Storage Policies to flush until the pool's aggregated locations go back below a configurable watermark. This critical check is done once per hour and spans Storage Pools.

Distribution of Data to Physical Storage Locations

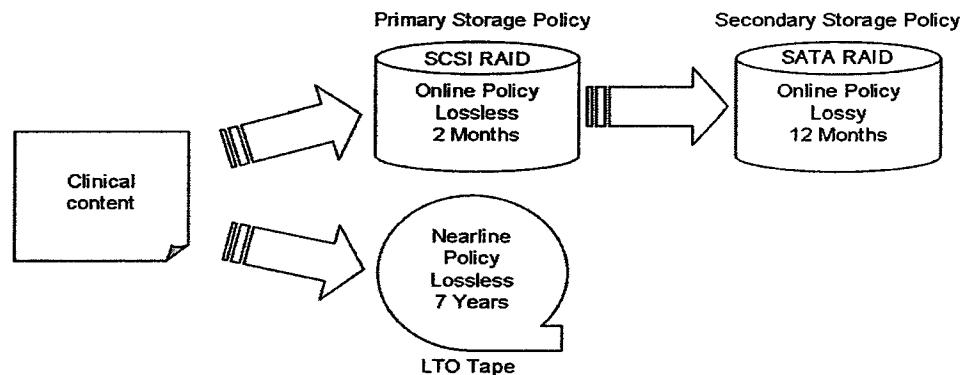
Content distribution within a storage pool to its locations is performed in a randomized manner. The pool randomly picks one of its locations that contain enough disk space for the content. If the incoming data belong to an existing study within the system, the new data are kept in the same location as the existing data. If the existing data is on a location that does not have enough space for the incoming data, the study is spread across locations, but within the same storage pool.



Life cycle management is used to minimize the cost of storage relative to the likelihood of content retrieval

Lifecycle Management Policies

There are additional lifecycle management rules that can be used to augment the aforementioned storage policy flushing rules. These lifecycle management rules allow content to be transferred to one or more other storage policies when it is flushed, as opposed to deleting it. Storage policies can be cascaded to enable true multi-tiered lifecycle management and is commonly used to minimize the cost of storage relative to the likelihood of content retrieval. This powerful feature can also be utilized to migrate content to new storage media transparently to the end users.



Clinical Lifecycle Management (example)

Special flushing rules allow pediatric studies to be retained in the system for longer periods of time to satisfy regulatory concerns.

Pediatric Policies

The Evercore flushing configuration also includes the option to enforce special flushing rules for pediatric studies. These rules allow pediatric studies to be retained in the system for longer periods of time to satisfy regulatory concerns. The standard Evercore flushing rules ignore pediatric studies that have not met the separate pediatric criteria for flushing.

UNIVISION - EMR/EHR INTEGRATION

Evercore provides a robust entry point for users throughout the enterprise to access up-to-date digital clinical content via their EMR/EHR.

Overview

It is not uncommon today to find that medical image and other digital clinical data necessary for quality patient care is locked away in departmental silos. Providing access to this data is often done using expensive PACS workstations or departmental PACS web viewers, if it is provided at all. Gaining access to a PACS workstation is often difficult because of cost and licensing issues. It can be also be very difficult (sometimes even impossible) to enable access to PACS viewing software in a managed environment, such as Citrix.

Because many providers utilize several different PACS for various imaging departments, EMR/EHR integration can also pose a challenge. Multiple HL7 interfaces must often be created and maintained to provide updates to the EMR/EHR from the departmental systems and it not uncommon for each PACS to utilize a different viewing application or technology. Thus, users must be trained to use multiple image viewers and are sometimes required to log in multiple times.

Evercore UniVision was designed to address these and other issues:

- The desire to have a common, lightweight (no downloaded components), browser-based image viewer, embedded into the EMR/EHR.
- The need to expose digital clinical content to a wider variety of users throughout the entire healthcare spectrum.
- The ability to preserve the look and feel and security attributes of the EMR/EHR that UniVision is embedded into.
- A single easy-to-use viewer for digital clinical content from multiple departments and in multiple formats (not just images).

Evercore UniVision is a robust entry point for users throughout the enterprise to access up-to-date digital clinical content via their EMR/EHR.

Evercore UniVision

UniVision suits the needs of general users by delivering an easy-to-use interface based on a Web browser using Web standards.

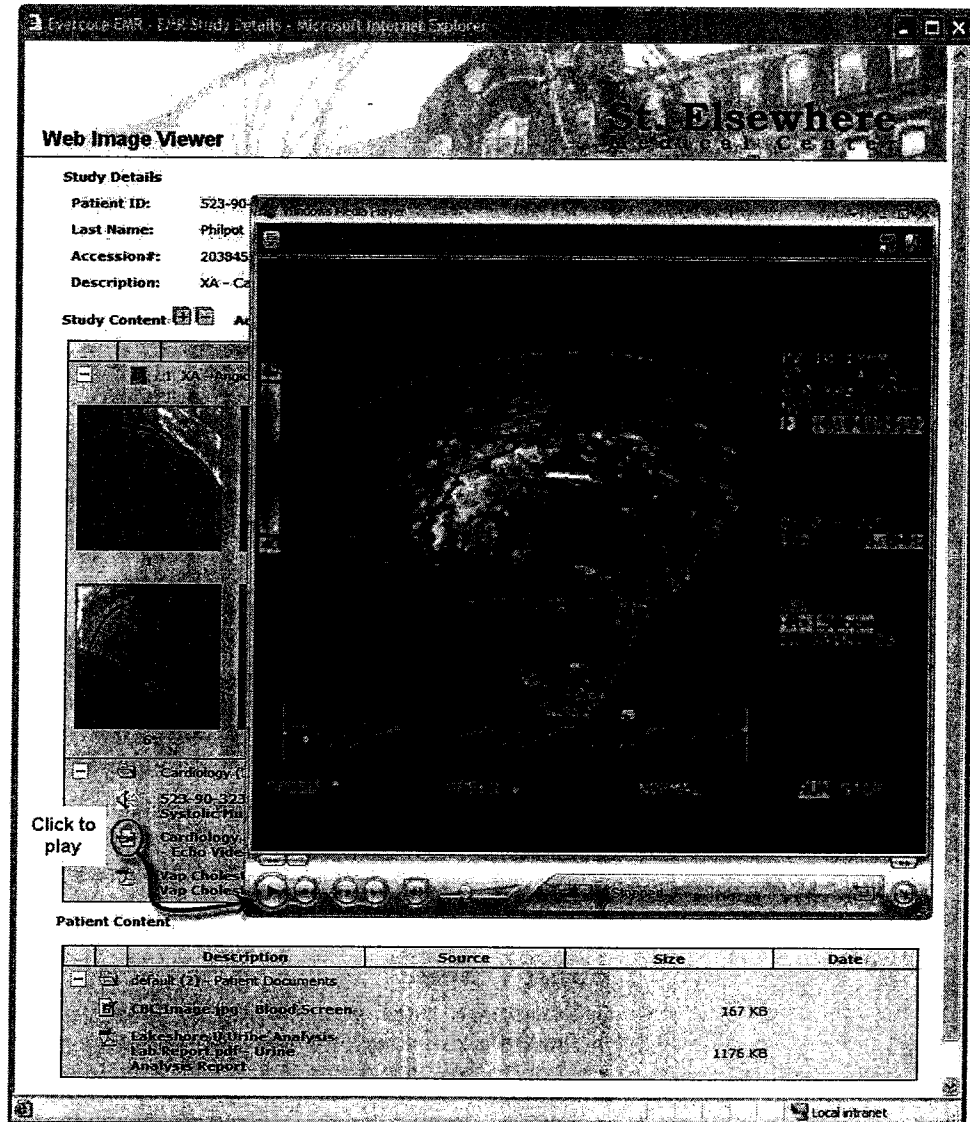
UniVision provides a simple solution to a complex problem and is especially relevant for referring physicians, general practitioners and nursing staff. These types of users often do not need the full fidelity images that PACS was designed to deliver. UniVision suits the needs of general users especially well by delivering an easy-to-use interface, completely through a Web browser using Web standards (HTML, XHTML, CSS, JPEG, etc).

When accessed via the EMR/EHR, UniVision immediately opens to the patient study context contained in the EMR/EHR. The simple user interface contains basic image study information, image thumbnails, and the ability to scroll through the images. An image can then be selected to view series and image information, along with an enlarged view of the image itself. While the images presented are slightly reduced in quality, they are sufficient for consultative purposes with patients or other staff.

UniVision supports the option to display diagnostic reports along with the images and other clinical content. Reports in both the HL7 and DICOM SR standard are supported. In addition to standard digital images and reports, users can optionally access non-DICOM clinical content that has been attached to the patient record or the patient's study. Users can view PDF files, video files, JPEG, and TIFF images as well as listen to sound files or access any other stored clinical data.

Using industry standards, UniVision can be delivered anywhere with network connectivity to the Evercore web server. With Evercore as the enterprise content manager and repository, EMR/EHR applications need only worry about one URL format to one enterprise system. This not only reduces interfacing costs, but also provides access to a wide variety of content from different clinical areas through a single, easy-to-use interface.

In addition to digital images, users can access non-DICOM clinical content that has been attached to the patient's study—PDF files, video files, sound files, and so forth.



UniVision EMR/EHR Viewer – Non-DICOM Video File

Study - Series Information

As you can see in the above picture, the study information is displayed along with thumbnail views of all series in the study. If a DICOM key image note exists, those images will be presented as the first series. If configured, the user can also access other studies which may exist for the patient. Most of the information displayed is easily customizable via CSS (cascading style sheets).

Users can manipulate a slider bar to scroll through the images. Clicking a thumbnail expands that image to a full-size JPEG which the user can print out

if desired. Users can also set up a slide show to view the images in sequence automatically.

Non-DICOM Clinical Content

Folders containing non-DICOM clinical content display below the thumbnail images. Non-DICOM content at both the study and/or patient level can be easily accessed (non-DICOM access is a configurable option). Users can click the file icon to view the contents in the browser's supported application.

SECURITY

Overview

In an era of increasing regulations surrounding access to Protected Health Information (PHI), the Evercore solution can facilitate enforcement of institutional security policies and procedures. Evercore maintains audit logs around every important aspect of data access and management within its system. The solution limits access to protected health data and logs system activities to enable better HIPAA compliance and monitoring.

HIPAA

TeraMedica takes HIPAA compliance very seriously and provides a flexible way for healthcare providers to both appropriately limit access to data while expanding the longitudinal clinical data record.

SmartStore Security

The organizational structure of Evercore SmartStore provides a mechanism to logically (and physically if necessary) segment the clinical data according to business unit needs. This segmentation creates a natural means to limit access to parts of the patient record based on the user's authority level and access point in the configured organization hierarchy.


The Evercore system implements user ID and password security for controlling authentication and access to the Evercore user interface. In addition, users can be assigned to groups which control authorization to various system features and functionality.

Audit Trails

Evercore maintains an audit trail for many system-wide activities and facilitates a provider's data security policies and procedures. These logs, which are kept within the database, can assist compliance to HIPAA regulations by documenting access, attempted access, and modification to protected health information. A web-based query tool searches the audit trail using a variety of criteria, including patient MRN and user id.

In addition, email notification of important audit-related system events can be configured to notify appropriate personnel. Auditing is performed for many Evercore user interface activities and selective auditing can be configured. Audit information for DICOM and HL7 interfaces is limited to data supplied in the processed request.

The audit trails are stored in the Evercore database according to existing IHE and DICOM standards for content.



For UniVision access via an EMH/EHR, users are authenticated within the EMR/EHR and do not need to login again to use UniVision, thus maintaining a more seamless integration.

Enterprise Security Integration

Evercore provides the ability to integrate with an existing LDAP-based security infrastructure (such as Active Directory Services) which provides users a single authentication context. Evercore can also be made to interface to an enterprise single sign-on application.

For UniVision access via an EMH/EHR, users are authenticated within the EMR/EHR and do not need to login again to use UniVision, thus maintaining a more seamless integration.

DEPLOYMENT

The Evercore solution can be deployed in a variety of configurations to accommodate existing technology standards and practices as well as enabling a scalable, robust infrastructure.

Overview

TeraMedica recognizes that most healthcare providers have ongoing relationships with their IT technology providers. The Evercore solution can be deployed in a variety of configurations to accommodate existing technology standards and practices as well as to enable a scalable, robust infrastructure.

Scalability

Evercore has been successfully deployed into very high-volume production use at healthcare systems where it is managing in excess of 2 million studies annually and more than 250 million clinical objects in total.

Flexibility in its architecture allows Evercore to be scaled either horizontally (more smaller servers) or vertically (fewer larger servers) to achieve a balance between data I/O, processing and costs.

Deployment Options

Server/Operating System

Evercore can be deployed on the following platforms:

- Sun Microsystems Solaris
- IBM AIX
- Microsoft Windows

Note: In a distributed deployment using Evercore FastGrid, platforms can be mixed or matched.

The solution processes can be provisioned in a variety of ways across single or multiple instances of these platforms to enable a highly scalable and available infrastructure.

Evercore SmartStore allows users to store their data on the types of media they desire with the rules and policies that are right for their part of the clinical or research practice.

Storage

Evercore supports a wide variety of storage subsystem options including:

- Direct Attached Storage (DAS)
- Network Attached Storage (NAS)
- Storage Area Networks (SAN)
- Content Addressed Storage (CAS)
- Grid Storage

Note: One or more of these types of storage can be used simultaneously with a single deployment of Evercore.

Additionally Evercore has been certified to work with the following Hierarchical Storage Managers (HSM):

- IBM Tivoli Storage Manager
- Sun Microsystems SAM/FS

These storage managers provide access to hundreds of different vendor storage solutions utilizing spinning disk, CD, DVD, tape, etc.

Evercore SmartStore allows users to store their data on the types of media they desire with the rules and policies that are right for their part of the clinical or research practice. One deployment of Evercore can utilize different storage types and formats for different users providing unique qualities of service for each business unit. In addition, new types of storage can be added dynamically to the system and existing data migrated to the new storage transparently to the end users.

Database

The following relational database management systems are supported for Evercore deployments:

- Oracle Enterprise
- Microsoft SQL Server
- PostgreSQL


Note: In a distributed deployment using Evercore FastGrid, databases can be mixed or matched.

FastGrid: Evercore Distributed Deployment

Overview

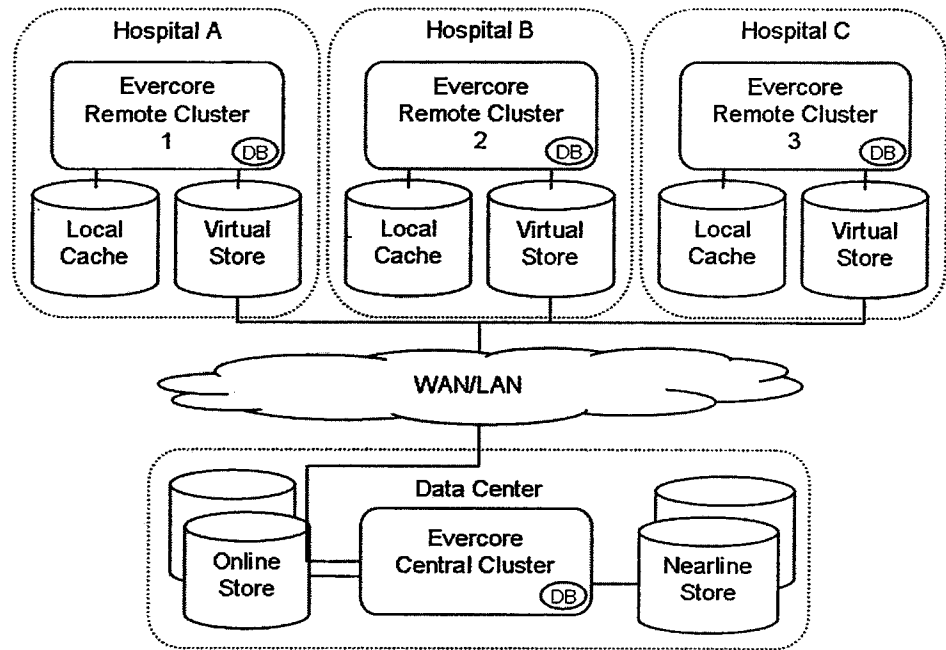
FastGrid extends Evercore out of a single cluster scenario into a true distributed environment with multiple clusters. This can be utilized for a geographically dispersed provider network, a group of hospitals or clinics that do not have high speed connections between them, or outlying imaging locations that desire to keep a local cache of clinical content but still remain connected to a central data center.

A cluster refers to one or more servers running Evercore (or portions of Evercore) which are connected to a single database instance. FastGrid enables remote sites to have an instance of Evercore on one or more servers. Consequently, remote sites have a local Evercore cache of data and access to metadata in the local database. This local cache can provide significant performance gains in terms of viewing, querying, and retrieving data at the remote site while improving business continuance.



FastGrid can be utilized for a geographically dispersed provider network, a group of hospitals or clinics that do not have high-speed connections between them, or outlying imaging locations that desire to keep a local cache of clinical content but still remain connected to a central data center.

The diagram below outlines the basic FastGrid concept:



Evercore FastGrid Deployment (example)

All image and content data transfers between Evercore clusters are conducted via an efficient, encrypted, recoverable communications protocol.

Communications

In addition to the local cache of data, the remote and central Evercore clusters are connected via two different communication pipes. All image and content data transfers between these clusters are conducted via an efficient, encrypted, recoverable communications protocol. Modifications to content metadata in the Evercore database (at either the remote or central site) are communicated using a bidirectional secure XML-based messaging scheme.

Users at the remote site interact with their local Evercore instance as usual, but in the background Evercore transparently retrieves the data that it does not have at that site. If the network connection between the sites is unavailable, the remote site still has its data cache to work from in order to provide a level of autonomy and business continuance.

Workflow

In order to provide its remote functionality, Evercore extends content workflow operations to reach across the FastGrid network. If a query is performed at a remote cluster, results are returned from that site's local database. If content is retrieved, the remote cluster knows if it is stored locally or offsite at the central cluster. If the content is offsite, a retrieve job is created in order to retrieve the data from the central cluster.

When data storage occurs at the remote cluster, the content can be cached at the remote site for rapid access. Additionally, the content can be transferred in the background to the central cluster for long-term storage and disaster backup. This transfer can be configured to occur in real time or in a batch mode during off-peak hours. If content metadata (patient name, MRN, etc.) is modified at either cluster, FastGrid sends updates to all relevant clusters in the FastGrid network.

If content metadata (patient name, MRN, etc.) is modified at either cluster, FastGrid sends updates to all relevant clusters in the FastGrid network.

These transactions and more are enabled by extending the SmartStore organization mapping definition to include the remote organizations. Therefore, an organization can exist at two clusters, each with its own content storage and management policies.

GLOSSARY

Term	Description
AE	Application Entity (AE) is a DICOM conformant agent on the network
AE Title	The name of an Application Entity. The name consists of printable ASCII characters and is limited to a length of 16. The AE Title is not the same as the host name.
ASCII	American Standard Code for Information Interchange (ASCII) is a code that represents letters, numerals, punctuation marks and control signals as seven bit groups. It is used as a standard code by the transmission of data.
CSS	Cascading Style Sheets (CSS), a new feature being added to HTML that gives both Web site developers and users more control over how pages are displayed. With CSS, designers and users can create style sheets that define how different elements, such as headers and links, appear. These style sheets can then be applied to any Web page.
DAS	Direct-Attached Storage (DAS) is computer storage that is directly attached to one computer or server and is not, without special support, directly accessible to other ones.
DICOM	Digital Imaging and Communications in Medicine (DICOM) is a healthcare standard developed by the American College of Radiology Manufacturers Association to define the connectivity and communication protocols of medical imaging devices.
DVD	Digital Video Disk (DVD) is an optical disc technology with a 4.7 gigabyte storage capacity on a single-sided, one-layered disk, which is enough for a 133-minute movie. DVDs can be single- or double-sided, and can have two layers on each side; a double-sided, two-layered DVD will hold up to 17 gigabytes of video, audio, or other information. This compares to 650 megabytes (.65 gigabyte) of storage for a CD-ROM disk.
EHR	Electronic Health Record (EHR) is an electronic patient record that resides in a system specifically designed to support users by providing accessibility to complete and accurate data, alerts, reminders, clinical decision support systems, links to medical knowledge, and other aids.
EMR	Electronic Medical Record (EMR) is the current term used to refer to computerization of health record content and associated processes within a healthcare enterprise.
HL7	An acronym for Health Level 7, it is a standard for healthcare and is the interface standard for communication between various systems employed in the medical community.
HTML	Hypertext Markup Language (HTML) is the set of markup symbols or codes inserted in a file intended for display on a World Wide Web browser page. The markup tells the Web browser how to display a Web page's words and images for the user.
HTTP	Hyper Text Transfer Protocol (HTTP) is the actual communications protocol that enables Web browsing.
IT	Information Technology (IT) is the branch of engineering that deals with the use of computers and telecommunications to retrieve and store and transmit information
JPEG	Joint Photographic Experts Group (JPEG) is a lossless or lossy compression technique for monochrome or color images. This compression mechanism has been approved for medical use as part of the DICOM standard.
LDAP	LDAP (Lightweight Directory Access Protocol) is a software protocol for enabling anyone to locate organizations, individuals, and other resources such as files and devices in a network, whether on the public Internet or on a corporate intranet.
LTO	Linear Tape Open (LTO) is an open format technology that was developed jointly by HP, IBM and Certance (Seagate) to provide a clear and viable choice in an increasingly complex array of tape storage options. LTO technology is technology, which means that users will have multiple sources of product and media and the open nature of LTO technology also provides a means of enabling compatibility between different vendors' offerings.
PACS	Picture Archiving and Communication System (PACS) is a system that acquires, transmits, stores, retrieves, and displays digital images and related patient information from a variety of imaging sources and communicates the information over a network.
PDF	Portable Document Format (PDF) is a file format that has captured all the elements of a printed document as an electronic image that you can view, navigate, print, or forward to someone else.

Term	Description
RAID	RAID (Redundant Array of Independent Disks) is a way of storing the same data in different places (thus, redundantly) on multiple hard disks. By placing data on multiple disks, I/O (input/output) operations can overlap in a balanced way, improving performance.
RIS	A Radiology Information System (RIS) is used by radiology departments to store, manipulate and distribute patient radiological data and imagery. The system generally comprises of patient tracking and scheduling, result reporting and image tracking capabilities.
RMI	RMI (Remote Method Invocation) is a way that a programmer, using the Java programming language and development environment, can write object-oriented programming in which objects on different computers can interact in a distributed network. RMI is the Java version of what is generally known as a remote procedure call (RPC), but with the ability to pass one or more objects along with the request.
SAN	A Storage Area Network (SAN) is a specialized network that provides access to high performance and highly available storage subsystems using block storage protocols. The SAN is made up of specific devices, such as host bus adapters (HBAs) in the host servers, switches that help route storage traffic, and disk storage subsystems. The main characteristic of a SAN is that the storage subsystems are generally available to multiple hosts at the same time, which makes them scalable and flexible.
SATA	Often abbreviated <i>SATA</i> or <i>S-ATA</i> , an evolution of the Parallel ATA physical storage interface. Serial ATA is a serial link - a single cable with a minimum of four wires creates a point-to-point connection between devices.
SCP	A Service Class Provider (SCP) is a role on an Application Entity that corresponds to a Server in a Client/Server Pair.
SCSI	Small Computer System Interface (SCSI) is a parallel interface standard for attaching peripheral devices to computers. SCSI interfaces provide for faster data transmission rates (up to 80 megabytes per second) than standard serial and parallel ports.
SCU	Service Class User (SCU) is a role of an Application Entity that corresponds to a Client in a Client/Server Pair.
TIFF	Tag Image File Format (TIFF) is a common format for exchanging raster graphics (bitmap) images between application programs, including those used for scanner images. One of the most common graphic image formats, TIFF files are commonly used in desktop publishing, faxing, 3-D applications, and medical imaging applications.
URL	Uniform Resource Locator (URL) is the unique address for a file that is accessible on the Internet. A common way to get to a Web site is to enter the URL of its home page file in your Web browser's address line. However, any file within that Web site can also be specified with a URL. Such a file might be any Web (HTML) page other than the home page, an image file, or a program such as a common gateway interface application or Java applet. The URL contains the name of the protocol to be used to access the file resource, a domain name that identifies a specific computer on the Internet, and a pathname, a hierarchical description that specifies the location of a file in that computer.
XDS	Cross-Enterprise Document Sharing. The XDS profile provides a standards-based specification for managing the interchange of documents that care delivery organizations have decided to explicitly share.
XHTML	XHTML (Extensible Hypertext Markup Language) is "a reformulation of HTML 4.0 as an application of the Extensible Markup Language (XML)." HTML is the set of codes (that's the "markup language") that a writer puts into a document to make it displayable on the World Wide Web. HTML 4 is the current version of it. XML is a structured set of rules for how one might define any kind of data to be shared on the Web. It's called an "extensible" markup language because anyone can invent a particular set of markup for a particular purpose and as long as everyone uses it (the writer and an application program at the receiver's end), it can be adapted and used for many purposes - including, as it happens, describing the appearance of a Web page. That being the case, it seemed desirable to reframe HTML in terms of XML. The result is XHTML, a particular application of XML for "expressing" Web pages.
XML	XML (Extensible Markup Language) is a W3C initiative that allows information and services to be encoded with meaningful structure and semantics that computers and humans can understand. XML is great for information exchange, and can easily be extended to include user-specified and industry-specified tags

